REMARKS

Claims 1-35 are pending in this matter. Claims 1-5, 9, 13-14, 16-18, 29, 31-32, and 34-35 were rejected under 35 USC 102(e) as being anticipated by Rabinovich (U.S. 6,256,675). Claims 6-8, 19, 30 and 33 were rejected under 35 USC 103(a) as being unpatentable over Rabinovich. Claims 10-11, 15, 21-22 were rejected under 35 USC 103(a) as being unpatentable over Rabinovich further in view of Shah, et al. (U.S. 6,298,381). Claims 23-28 were rejected under 35 USC 103(a) as being unpatentable over Aggarwal et al. (U.S. 5,924,116) further in view of Rabinovich. The examiner objected to claim 12 as being dependent upon a rejected base claim but noted that it would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims. Applicant thanks the examiner for this indication of allowable matter.

Rejections Under 35 U.S.C. §§ 102 and 103

As to claim 1, the examiner indicated on page 2 of the office action that Rabinovich teaches determining distance metrics with respect to a first server or host, citing col. 7, lines 45-55. But Rabinovich's teaching as to the nature of the host and the requestor are distinguishable from the server and clients recited in claim 1. Rabinovich qualifies that distance metric as between a host p that stores a replica and the requestor. Such a host storing a replica is distinguishable from a server capable of storing and serving content. Moreover, Rabinovich's use of the term "requestor", such as exemplified by the requestor 109 shown in FIG. 1, cannot reasonably be categorized as dealing with anticipating need by clients. As Rabinovich further points out, the result of the analysis is the selection of a host to respond to the request (7:39-44). Thus, the requestor in Rabinovich isn't a client likely to access the content but an entity actually requesting the content. Hence, Rabinovich does not relate to determining either the location of clients that are likely to access the content or determining a first proximity between the client or group of clients and a first server capable of storing and serving the content. Instead Rabinovich relates to actual requests and servers already storing replicas. Accordingly, for at least the reasons cited above, applicants submit that Rabinovich fails to teach or suggest all of the elements of claim 1 including determining the location of a client or group of clients that are likely to access the content and determining a first proximity between the client or group of clients and a first server capable of storing and serving the content.

For the same reasons as discussed above with respect to claim 1, applicants submit that Rabinovich fails to teach or suggest all of the elements of claim 13.

Claim 14 has been amended to further indicate measures of proximities used. Support for the amendment may be found throughout the specification, including page 10, lines 20-30. Although Rabinovich notes that the distance metric measures the cost of communicating between the requestor and the host, the only examples provided are latency between the requestor and the host and, in a separate example, the bandwidth of the channel between the requestor and the host. Rabinovich makes no teaching or suggestion that the "cost" comprises any specific measures other than one of the two mentioned. Hence, for at least this reason, Rabinovich fails to teach or suggest all of the elements of claim 14.

As to claim 16, the examiner indicated that at col. 8, Rabinovich disclosed that the request for content is sent to the host that has the best fit distance, cost or delay metrics with respect to the requesting client. Applicant disagrees that Rabinovich teaches or suggests the limitations of claim 16 including wherein at least one of the first and second proximities is determined by at least one of the following factors: congestion, noise and loss on a network segment, and charges incurred to send. Although Rabinovich notes that the distance metric measures the cost of communicating between the requestor and the host, the only specific teachings provided are latency between the requestor and the host and, in a separate example, the bandwidth of the channel between the requestor and the host. Rabinovich makes no teaching or suggestion that the "cost" comprises any specific measures other than one of the two mentioned. Hence, for at least this reason, Rabinovich fails to teach or suggest all of the elements of claim 16.

Claims 23 -28 were rejected by the examiner under 35 USC 102(e) as being unpatentable over Aggarwal et al. (U.S. 5,924,116) further in view of Rabinovich. The examiner stated that Aggarwal teaches a method of releasing stored control items from a server to make room for new content items but acknowledged on page 7 of the office action that Aggarwal does not explicitly teach the limitation "wherein at least one of the first and second proximities is determined by a combination of the following factors: bandwidth, number of hops, congestion, noise and loss on a network segment, and charges incurred to send". The examiner further stated that Rabinovich teaches the limitation in a system and method for allocating requests for objects and managing replicas of objects. Applicant disagrees with the examiner's analysis of Aggarwal and Rabinovich.

Aggarwal deals with proxy servers acting as caches for Internet materials. Although Aggarwal generally discusses replacement selection logic, its tentative identification of objects for replacement is distinguishable from the recitations of claim 23. Specifically, Aggarwal does not teach or suggest a comparison of the relative values of the first and second proximities. Claim 23 has been amended to clarify the claimed invention, i.e., that "releasing one of the first and second stored content items based upon the relative values of the first and second proximities," involves releasing one of the first and second stored content items by directly comparing the relative values of the first and second proximities. Support for the amendment may be found in the specification at page 17, lines 15-28. In contrast, Aggarwal makes no comparison between proximities and their relative values but instead teaches selecting as the tentative object file to purge those objects having the largest product of time since last access and selection weights (col. 10, lines 20-25, 50-55). Even Aggarwal's selection weight fails to match applicant's description of proximities. Aggarwal describes the selection weight as a function of an object's size divided by the replacement cost of obtaining the object if it is not cached (col. 9, lines 52-55). Thus, Aggarwal makes no comparison of proximities and their relative values but performs several manipulations on the replacement cost of obtaining the object before determining which is a tentative candidate for replacement (See FIG. 7, step 550). Hence, Aggarwal fails to teach or suggest all of the limitations of claim 23 including releasing one of the first and second stored content items by directly comparing the relative values of the first and second proximities

Moreover, applicant disagrees that Rabinovich's teachings or suggestions as to distance metrics teach the specific limitations as recited in dependant claim 28. Although Rabinovich notes that the distance metric measures the cost of communicating between the requestor and the host, the only examples provided are latency between the requestor and the host and, in a separate example, the bandwidth of the channel between the requestor and the host. Rabinovich makes no teaching or suggestion that the "cost" comprises any specific measures other than one of the two mentioned and thus makes no suggestion as to using those two or any two in combination (as in dependant claim 28). Hence, applicant respectfully submits that Aggarwal and Rabinovich, either alone or in combination, fail to teach or suggest all of the elements of claim 23 and further, dependant claim 28.

Claim 29 is a content control system claim. Applicant disagrees that Rabinovich's teachings or suggestions as to distance metrics teach the specific limitations as recited in the

claim. Although Rabinovich notes that the distance metric measures the cost of communicating between the requestor and the host, the only examples provided are latency between the requestor and the host and, in a separate example, the bandwidth of the channel between the requestor and the host. Rabinovich makes no teaching or suggestion that the "cost" comprises any specific measures other than one of the two mentioned. Hence, applicant respectfully submits that Rabinovich fails to teach or suggest all of the elements of claim 29.

Claims 2-12, 15, 17-22, 24-28, and 30-35 are dependant claims, depending respectively form independent claims 1, 14, 16, 23, and 29. Thus, at least due to these dependencies, the dependent claims are submitted to be in allowable form. Further, the dependent claims recite additional elements which when taken in the context of the claimed invention further patentably distinguish the art of record. The additional limitations recited in the dependent claims are not further discussed as the above-discussed limitations are clearly sufficient to distinguish the claimed invention from Aggarwal and Rabinovich. Withdrawal of the rejections is respectfully requested.

Conclusion

Accordingly, it is submitted that all issues in the Office Action have been addressed, and withdrawal of the rejections is respectfully requested. Applicant believes that this application is in condition for allowance, and requests a prompt passage to issuance. If the Examiner believes that a telephone conference would expedite the prosecution of this application, he is invited to contact the Applicant's undersigned attorney at the telephone number set out below.

Respectfully submitted,
BEYER WEAVER & THOMAS, LLP

Russell N. Swerdon
Registration No. 36,943

P.O. Box 778
Berkeley, CA 94704-0778
(510) 843-6200

APPENDIX

VERSION WITH MARKINGS TO SHOW CHANGES MADE

CLAIMS

14. (once amended) A method of loading content to a server in anticipation of need by network clients, the method comprising:

(a) determining the location of a client or group of clients that require access to the content;

(b) identifying a first server currently storing the content and serving the client's requirements for access to the content;

(c) identifying a second server that does not currently store said content but that has the capability of storing and serving the content;

(d) determining a first proximity between the first server and the client or group of clients;

(e) determining a second proximity between the second server and the client or group of clients, wherein at least one of the first and second proximities is determined by at least one of the following factors: congestion, noise and loss on a network segment, and charges incurred to send ; and

(f) if the relative values of the first and second proximities meet a defined constraint, loading the content to the second server.

23. (twice amended) A method of releasing stored content items from a server to make room for new content items, the method comprising:

(a) identifying, on the server, a first stored content item and a second stored content item;

(b) determining a first proximity between the server and a source of the first stored content item;

(c) determining a second proximity between the server and a source of the second stored content item; and

(d) releasing one of the first and second stored content items by directly comparing [based upon] the relative values of the first and second proximities, wherein at least one of the first and second proximities is determined by at least one of the following factors: bandwidth, number of hops, congestion, noise and loss on a network segment, and charges incurred to send.

APPENDIX OF PENDING CLAIMS

1. A method of loading content to a server in anticipation of a need for the content by network clients, the method comprising:

(a) determining the location of a client or group of clients that are likely to access the content;

(b) determining a first proximity between the client or group of clients and a first server capable of storing and serving the content;

(c) determining a second proximity between the client or group of clients and a second server capable of storing and serving the content; and

(d) based upon the relative values of the first and second proximities, loading the content into one of the first and second servers.

2. The method of claim 1, wherein loading the content to the second server is performed automatically by a content control system on the network.

3. The method of claim 2, wherein performing (b), (c), and (d) is accomplished by the content control system.

4. The method of claim 1, wherein the first and second proximities are determined dynamically by a content control system.

5. The method of claim 1, wherein the content is loaded to the server that is most proximate the client or group of clients.

6. The method of claim 1, wherein the content is multimedia content.

7. The method of claim 6, wherein the multimedia content is transmitted over the network in a compressed format.

8. The method of claim 1, wherein the content is video content.

9. The method of claim 1, wherein at least one of the first and second proximities is determined by a combination of the following factors: bandwidth, number of hops, congestion, noise and loss on a network segment, and charges incurred to send.

10. The method of claim 1, wherein at least one of the first and second proximities is determined by a considering whether the server and the client or group of clients are on the same sub-network.

11. The method of claim 10, wherein content is loaded to the second server when the second server and the client or clients are on the same sub-network and the first server and the client or clients are not on the same sub-network.

12. The method of claim 1, further comprising:

determining a first loading proximity between a source of the content and the first server;

determining a second loading proximity between a source of the content and the second server; and

using the first and second loading proximities together with the first and second proximities to determine which of the first and second servers should receive the content.

13. A method of loading content to a server in anticipation of a need for the content by network clients, the method comprising:

(a) determining the location of a client or group of clients that are likely to access the content;

(b) determining the relative proximities of (i) the client or group of clients to a first server capable of storing and serving the content and (ii) the client or group of clients and a second server capable of storing and serving the content; and

(c) based upon the relative proximity, loading the content into one of the first server and the second server.

14. (once amended) A method of loading content to a server in anticipation of need by network clients, the method comprising:

(a) determining the location of a client or group of clients that require access to the content;

(b) identifying a first server currently storing the content and serving the client's requirements for access to the content;

(c) identifying a second server that does not currently store said content but that has the capability of storing and serving the content;

(d) determining a first proximity between the first server and the client or group of clients;

(e) determining a second proximity between the second server and the client or group of clients, wherein at least one of the first and second proximities is determined by at least one of the following factors: congestion, noise and loss on a network segment, and charges incurred to send ; and

(f) if the relative values of the first and second proximities meet a defined constraint, loading the content to the second server.

15. The method of claim 14, wherein (d) includes determining whether the server and the client or group of clients are on the same sub-network.

16. (once amended) A method of selecting a server to fill a client request for content, the method comprising:

(a) determining that one or more clients needs or will need to receive the content;

(b) determining a first proximity between the one or more clients and a first server capable of supplying the content;

(c) determining a second proximity between the one or more clients and a second server capable of supplying the content; and

(d) based upon the relative values of the first and second proximities, choosing one of the first and second servers to fill client requests for the content, wherein at least one of the first and second proximities is determined by at least one of the following factors: congestion, noise and loss on a network segment, and charges incurred to send.

17. The method of claim 16, wherein the first and second proximities are determined dynamically by a content control system.

18. The method of claim 16, wherein the content is provided by the server that is most proximate to the one or more of clients.

19. The method of claim 16, wherein the content is multimedia content.

20. The method of claim 16, wherein at least one of the first and second proximities is determined by a combination of the following factors: bandwidth, number of hops, congestion, noise and loss on a network segment, and charges incurred to send.

21. The method of claim 16, wherein the first proximity is determined by a considering whether the first server and the one or more clients are on the same sub-network.

22. The method of claim 21, wherein content is provided by the first server when the first server and the one or more clients are on the same sub-network and the second server and the one or more clients are not on the same sub-network.

23. (twice amended) A method of releasing stored content items from a server to make room for new content items, the method comprising:

(a) identifying, on the server, a first stored content item and a second stored content item;

(b) determining a first proximity between the server and a source of the first stored content item;

(c) determining a second proximity between the server and a source of the second stored content item; and

(d) releasing one of the first and second stored content items by directly comparing the relative values of the first and second proximities, wherein at least one of the first and second proximities is determined by at least one of the following factors: bandwidth, number of hops, congestion, noise and loss on a network segment, and charges incurred to send.

24. The method of claim 23, wherein at least one of the sources of the first and second stored content items is a content library.

25. The method of claim 23, wherein the stored content items are video content items.

26. The method of claim 23, wherein the first and second stored content items are identified based upon a cache release protocol.

27. The method of claim 26, wherein the cache release protocol is a Least Recently Used algorithm.

28. The method of claim 23, wherein at least one of the first and second proximities is determined by a combination of the following factors: bandwidth, number of hops, congestion, noise and loss on a network segment, and charges incurred to send.

29. (once amended) A content control system for propagating content on a network, the content control system comprising:

an interface to the network; and

a processor and a memory coupled to said processor, the processor and memory configured or designed to determine proximities of network nodes to one another and to propagate content to one of said nodes based upon a proximity determination, wherein at least one of the proximities is determined by at least one of the following factors: congestion, noise and loss on a network segment, and charges incurred to send.

30. The content control system of claim 29, wherein the interface, processor, and memory are provided on a router.

31. The content control system of claim 29, wherein the interface, processor, and memory are provided on a PC or workstation.

32. The content control system of claim 29, further comprising an operating system.

33. The content control system of claim 29, further comprising a video server running on the operating system.

34. The content control system of claim 29, further comprising a mass storage device capable of storing content and a mass storage controller capable of controlling access to content stored in the mass storage device.

35. The content control system of claim 29, further comprising one or more proximity determining tools.